REMARKS

Claims 1-30 are pending. Claims 10-26 were previously withdrawn. No new matter has been added.

Claims 1-6 were rejected as being unpatentable over the combination of Lum (U.S. Patent No. 5,696,931, hereinafter "Lum") over Hicken (U.S. Patent No. 6,092,149, hereinafter "Hicken") in view of Napolitano (U.S. Patent No. 6,301,605, hereinafter "Napolitano"). Applicant respectfully traverses the section 103(a) rejection because Lum, Hicken, and Napolitano do not teach or suggest all of the limitations of the claims as currently provided.

Claim 1 requires that the controller "requests a transfer of the requested data that resides in the mass storage device directly to the host system." Examiner, in the Office Action dated August 22, 2007, has conceded that neither Lum nor Hicken disclose this claim element (see Office Action paragraph bridging pages 4 and 5). Examiner has asserted, however, that Napolitano discloses the direct transfer of the requested data from the mass storage device to the host system. Applicant respectfully traverses that assertion.

Napolitano discloses in column 10, lines 19-35 that "The file array system 410 also collects statistics to exploit the distributed nature of the architecture. These statistics generally relate to (i) sequentiality, (ii) cache hit rates, (iii) contention for determining read-ahead aggressiveness, (iv) cache thrashing, and (v) bypassing of adapter cache...**the server file system uses these statistics to determine whether to send information to the client file system before the client requests the information** (*emphasis added*). For example, the server file system **may determine that there is bandwidth available** (*emphasis added*) on the SCSI

channels and that it is likely that **the client file system may request certain information** (*emphasis added*). Accordingly, the DMA engine 356 **transfers the information** (*emphasis added*) obtained from disks directly into host memory 330 so that it is available at the host **if and when it is requested** (*emphasis added*)."

Rather than transferring the **requested data** from the mass storage device to the host system, Napolitano discloses the predictive retrieval of **unrequested data** ("...it is likely that the client file system may request certain information...transfers the information...so that it is available at the host if and when it is requested") into host memory 330.

Furthermore, in column 10, lines 62-65, Napolitano discloses "If the requested data is not in host cache 340 (Step 506), the read request is forwarded by the client file system 420 over the interface 302 to the server file system 450 in Step 512." Also, in column 11, lines 9-14, Napolitano discloses "If the data is not present in the adapter cache (Step 516), the server file system 450 issues an I/O data transfer request to the I/O subsystem 480 which retrieves the data from disk 325 in Step 522. **Once retrieved, the data is loaded into adapter memory 370** (*emphasis added*) and the process returns to Step 518 where **the DMA engine 356 transfers the requested data directly into host memory 330 via a DMA operation** (*emphasis added*)."

As is clear from the above quotations, rather than transferring the data from the mass storage device to the host system, Napolitano discloses that the data is transferred into adapter memory 370 and then a DMA engine 356 transfers the data into host memory 330. Again, this is contrary to Applicant's claim element of transferring "the requested data that resides in the mass storage device directly to the host system."

Additionally, in column 11, lines 21-31, Napolitano discloses "...such parallel execution not only "frees" the host CPU 312 for other processing operations, ...facilitates scaling of host

CPUs ...the burden of transaction implementation onto the adapter, the number of host CPU interrupts is reduced; ...the file system need only interrupt the host CPU 312 once when a transaction completes. By limiting the frequency of CPU interruption ...the file array architecture provides a substantial data processing performance enhancement." Napolitano discloses that the use of the host CPU 312 to perform transactions can negatively impact data processing performance. In other words, Napolitano teaches away from the claim limitation of transferring requested data directly to the host system.

Accordingly, Claim 1 is patentable over Lum, Hicken, and Napolitano, and respectfully request withdrawal of the Examiner's rejection.

Claims 2-6 depend from Claim 1. Applicant respectfully submits that these dependent claims are patentable over the cited prior art, not only because of their dependency from Claim 1 for the reasons discussed above, but also in view of their novel claim features.

Claims 8, 27-28, and 30 were rejected as being unpatentable over the combination of Lum over Simionescu (U.S. Patent No. 6,141,728, hereinafter "Simionescu") in view of Napolitano. Applicant respectfully traverses the section 103(a) rejection because Lum, Simionescu, and Napolitano do not teach or suggest all of the limitations of the claims as currently provided.

Claim 8 requires that "if a portion of the requested data is in the cache memory and a portion of the requested data is in the mass storage device, transferring the portion of the requested data from the cache memory to the host system substantially concurrently with transferring the portion of the requested data from the mass storage device directly to the host system." And Claim 27-28 require that "to cause said microprocessor to fetch data-blocks

corresponding to said cache-miss-portion of said data-request directly from said disk-device." Examiner has conceded that neither Lum nor Simionescu disclose these claim limitations (see Office Action dated August 22, 2007, at p. 8, first full paragraph, and p. 10, first full paragraph). Examiner has asserted, however, that Napolitano discloses the direct transfer of the requested data from the mass storage device to the host system. Applicant respectfully traverses that assertion.

Napolitano discloses in column 10, lines 19-35 that "The file array system 410 also collects statistics to exploit the distributed nature of the architecture. These statistics generally relate to (i) sequentiality, (ii) cache hit rates, (iii) contention for determining read-ahead aggressiveness, (iv) cache thrashing, and (v) bypassing of adapter cache...**the server file system uses these statistics to determine whether to send information to the client file system before the client requests the information** (*emphasis added*). For example, the server file system **may determine that there is bandwidth available** (*emphasis added*) on the SCSI channels and that it is likely that **the client file system may request certain information** (*emphasis added*). Accordingly, the DMA engine 356 **transfers the information** (*emphasis added*) obtained from disks directly into host memory 330 so that it is available at the host **if and when it is requested** (*emphasis added*)."

As is clear from the above quotations, rather than transferring the requested data from the mass storage device to the host system, Napolitano discloses the predictive retrieval of unrequested data ("...it is likely that the client file system may request certain information. Accordingly, the DMA engine 356 transfers the information obtained from disks directly into host memory 330 so that it is available at the host if and when it is requested") into host memory

330. This is contrary to Applicant's claim element of transferring 'the requested data that resides in the mass storage device directly to the host system."

Furthermore, in column 10, lines 62-65, Napolitano discloses "If the requested data is not in host cache 340 (Step 506), the read request is forwarded by the client file system 420 over the interface 302 to the server file system 450 in Step 512." Also, in column 11, lines 9-14, Napolitano discloses "If the data is not present in the adapter cache (Step 516), the server file system 450 issues an I/O data transfer request to the I/O subsystem 480 which retrieves the data from disk 325 in Step 522. **Once retrieved, the data is loaded into adapter memory 370** (*emphasis added*) and the process returns to Step 518 where **the DMA engine 356 transfers the requested data directly into host memory 330 via a DMA operation** (*emphasis added*)."

Clearly, rather than transferring the data from the mass storage device to the host system, Napolitano discloses that the data is transferred into adapter memory 370 and then a DMA engine 356 transfers the data into host memory 330.

Additionally, in column 11, lines 21-31, Napolitano discloses "...such parallel execution not only "frees" the host CPU 312 for other processing operations, ...facilitates scaling of host CPUs ...the burden of transaction implementation onto the adapter, the number of host CPU interrupts is reduced; ...the file system need only interrupt the host CPU 312 once when a transaction completes. By limiting the frequency of CPU interruption ...the file array architecture provides a substantial data processing performance enhancement." Napolitano discloses that the use of the host CPU 312 to perform transactions can negatively impact data processing performance. In other words, Napolitano teaches away from the claim limitation of transferring requested data directly to the host system.

Accordingly, Applicant submits that Claim 8 and Claim 27-28 are patentable over Lum, Simionescu, and Napolitano, and respectfully request withdrawal of the Examiner's rejection.

Claims 30 depends from Claim 28. Applicant respectfully submits that this dependent claim is patentable over the cited prior art, not only because of their dependency from Claim 28 for the reasons discussed above, but also in view of their novel claim features.

Claim 7 was rejected as being unpatentable over the combination of Lum over Hicken in view of Napolitano and well-known practices in the art. Applicant respectfully traverses the section 103(a) rejection because Lum, Hicken, Napolitano and well-known practices in the art do not teach or suggest all of the limitations of the claims as currently provided.

Claim 1 requires that that the controller "requests a transfer of the requested data that resides in the mass storage device directly to the host system." As discussed at length above, neither Lum, Hicken, Napolitano nor well-known practices in the art, taken alone or in combination, discloses the direct transfer of requested data residing in the mass storage device directly to the host system.

Claim 7 depends from Claim 1. Applicant respectfully submits that this dependent claim is patentable over the cited prior art, not only because of its dependency from Claim 1 for the reasons discussed above, but also in view of its novel claim features.

Claim 29 was rejected as being unpatentable over the combination of Lum over Simionescu in view of Napolitano and well-known practices in the art. Applicant respectfully traverses the section 103(a) rejection because Lum, Simionescu, Napolitano and well-known

practices in the art do not teach or suggest all of the limitations of the claims as currently provided.

Claim 28 requires that "to cause said microprocessor to fetch data-blocks corresponding to said cache-miss-portion of said data-request directly from said disk-device." Neither Lum, Simionescu, Napolitano nor well-known practices in the art, taken alone or in combination, discloses the direct transfer of requested data residing in the mass storage device directly to the host system.

Claim 29 depends from Claim 28. Applicant respectfully submits that this dependent claim is patentable over the cited prior art, not only because of its dependency from Claim 28 for the reasons discussed above, but also in view of its novel claim features.

Claim 9 was rejected as being unpatentable over the combination of Lum over Simionescu in view of Napolitano and U.S. Patent Publication No. 2001/0014929 to Taroda et al. (hereinafter, "Taroda") and well-known practices in the art. Applicant respectfully traverses the section 103(a) rejection because Lum, Simionescu, Napolitano, Taroda and well-known practices in the art do not teach or suggest all of the limitations of the claims as currently provided.

Claim 8 requires that "if a portion of the requested data is in the cache memory and a portion of the requested data is in the mass storage device, transferring the portion of the requested data from the cache memory to the host system substantially concurrently with transferring the portion of the requested data from the mass storage device directly to the host system." Neither Lum, Simionescu, Napolitano, Taroda, nor well-known practices in the art, taken alone or in combination, discloses the direct transfer of requested data residing in the mass storage device directly to the host system.

Claim 9 depends from Claim 8. Applicants respectfully submit that this dependent claim is patentable over the cited prior art, not only because of its dependency from Claim 8 for the reasons discussed above, but also in view of its novel claim features.

In view of the above, Applicant respectfully submits that the application is in condition for allowance and request that the Examiner pass the case to issuance. If the Examiner should have any questions, Applicant requests that the Examiner please contact Applicant's attorney at the address below. No fee is believed due in connection with this filing. However, in the event that there are any fees due, please charge the same, or credit any overpayment, to Deposit Account No. 50-1065.

Respectfully submitted,

November 14, 2007

Steven H. Slater
Reg. No. 35,361
Attorney for Applicant

SLATER & MATSIL, L.L.P.
17950 Preston Rd.
Suite 1000
Dallas, Texas 75252
Tel.: 972-732-1001
Fax: 972-732-9218